This article is obsolete. Its contents are now covered in arXiv 1511.09319, where we discuss a comprehensive system for behavior discovery and spatial alignment of articulated object classes from unstructured video (available at https://arxiv.org/abs/1511.09319)

arXiv:1412.0477v3 [cs.CV] 16 Aug 2016

Recovering Spatiotemporal Correspondence between Deformable Objects by Exploiting Consistent Foreground Motion in Video

Luca Del Pero[1] ldelper@inf.ed.ac.uk
Susanna Ricco[2] ricco@google.com
Rahul Sukthankar[2] sukthankar@google.com
Vittorio Ferrari[1] ferrari@inf.ed.ac.uk
[1]University of Edinburgh [2]Google Research

Abstract

Given unstructured videos of deformable objects, we automatically recover spatiotemporal correspondences to map one object to another (such as animals in the wild). While traditional methods based on appearance fail in such challenging conditions, we exploit consistency in object motion between instances. Our approach discovers pairs of short video intervals where the object moves in a consistent manner and uses these candidates as seeds for spatial alignment. We model the spatial correspondence between the point trajectories on the object in one interval to those in the other using a time-varying Thin Plate Spline deformation model. On a large dataset of tiger and horse videos, our method automatically aligns thousands of pairs of frames to a high accuracy, and outperforms the popular SIFT Flow algorithm.

1. Introduction

Most computer vision systems cannot take advantage of the abundance of Internet video content as training data. This is because current algorithms typically learn under strong supervision and annotating video content is expensive. Our goal is to remove the need for expensive manual annotations and instead reliably recover spatiotemporal correspondences between deformable objects under weak supervision. For instance, given a collection of animal documentary videos, can we automatically match pixels on a tiger in one video to those on a different tiger in another video (Figs. 1 and 4)?

Recovering point-to-point spatiotemporal correspondences across videos is powerful because it enables to assemble a collection of *aligned* foreground masks from a collection of videos of the same object class (Fig. 1). Accomplishing this task in the presence of significant object appearance variations is particularly important in order to capture the richness of the visual concept (*e.g.*, different coloring and textures of an animal). Achieving this could replace the expensive manual annotations required by several popular methods for learning visual concepts [9, 10, 14, 19, 36, 39], including methods that require annotations at the part level [1, 3, 15]. Additionally, it can enable novel applications, such as replacing one instance of an object with a suitable instance from a different video (like the orange and the white tiger in Fig. 1).

Figure 1. We recover point-to-point spatiotemporal correspondences across a collection of unstructured videos of deformable objects. Here, we display the recovered correspondences by mapping tigers from frames in two different videos (top) onto each other (bottom). Our method maps each part of a tiger in the first video to the corresponding part in the second (*e.g.*, head goes to head, front-right paw goes to front-right paw, etc.). We use motion cues to find short video intervals where the foreground moves in a consistent manner. This enables finding correspondences despite large variations in appearance (*e.g.*, white and orange tigers).

Instances of the same class in different videos exhibit large variations in appearance. Hence, traditional methods for matching still images using local appearance descriptors [2, 21, 26, 27] typically do not find reliable correspondences. We do this more effectively by aligning short temporal intervals where the objects exhibit consistent motion patterns. We exploit the characteristic motion of an object

class (*e.g.*, a tiger's prowl) to identify suitable correspondences, and combine motion and edge features to align them with great accuracy.

We present a new technique to align two sequences of frames spatiotemporally using a set of Thin Plate Splines (TPS), an expressive non-rigid mapping that has primarily been used for registration [7] and shape matching [16] in still images. We extend these ideas to video by fitting a TPS that varies smoothly in time to minimize the distance between edge points in corresponding frames from the two sequences.

We evaluate our method on a new set of ground-truth annotations: 19 landmarks (*e.g.*, left eye, front left knee, neck, etc.) for two classes (horses and tigers). We annotated ∼100 video shots per class, for a total of ∼35,000 annotated frames (25 minutes of video). The tiger shots come from a dataset of high-quality nature documentaries filmed by professionals [11]. The horse shots are sourced from the YouTube-Objects dataset [29], which are primarily low-resolution footage filmed by amateurs. This enables quantitative analysis on a large scale in two different settings. Experiments show that our method recovers around a 1000 pairs of correctly aligned sequences from 100 real-world video shots of each class. As the recovered alignment is between *sequences*, this amounts to having correspondences between 10,000 pairs of frames. This significantly outperforms the traditional approach of matching SIFT keypoints [27] and the popular SIFT Flow algorithm [26].

The contributions of our work are: (1) a weakly supervised system that goes from a large collection of unstructured video of an object class to a tight network of spatiotemporal correspondences between object instances; (2) a method for aligning sequences of frames with consistent motion using TPS; (3) publicly releasing the ground-truth annotations above. To our knowledge, this is the largest benchmark for sequence alignment to date.

2. Related work

Still image alignment. Most works on spatial alignment focus on matching between images for a variety of applications such as multi-view reconstruction [33], image stitching [4], and object instance recognition [17, 27]. The traditional approach identifies candidate matches using a local appearance descriptor (*e.g.*, SIFT [27]) with global geometric verification performed using RANSAC [8, 18] or semi-local consistency checks [17, 23, 32]. PatchMatch [2] and SIFT Flow [26] generalize this notion to match patches between semantically similar scenes.

Sequence alignment. Our method differs from previous work on sequence alignment [5,6,35] in several ways. First, we find correspondences between different scenes, rather than between different views of the same scene [5,6]. While the method in [35] is able to align actions across different scenes by directly maximizing local space-time correlations, it cannot handle the large intra-class appearance variations and diverse camera motions present in our videos. As another key difference, all above approaches require temporally pre-segmented videos. Instead, we operate on unsegmented videos and our method automatically finds which portions of each can be spatiotemporally aligned. Finally, these works have been evaluated only qualitatively on 5-10 pairs of sequences, providing no quantitative analysis.

In the context of video action recognition, there has been work on matching of spatiotemporal templates to actor silhouettes [20, 40] or groupings of supervoxels [24]. Our work is different because we map pixels from one unstructured video to another. The method in [22] mines discriminative space-time patches and matches them across videos. It focuses on rough alignment using sparse matches (typically one patch per clip), while we seek a finer, non-rigid spatial alignment. Other works on sequence alignment focus on temporal rather than spatial alignment [31] or target a very specific application, like aligning presentation slides to videos of the corresponding lecture [13].

TPS alignment. TPS were developed as a general purpose smooth functional mapping for supervised learning [37]. TPS have been used for non-rigid point matching between still images [7], and to match shape models to images [16]. The computer graphics community recently proposed semi-automated video morphing using TPS [25]. However, this method requires manual point correspondences as input, and it matches image brightness directly.

Learning from videos. A few recent works exploit video as a source of training data for object class detectors [29, 34]. However, their use of video is limited to segmenting objects from their background. Ramanan *et al*. [30] build a simple 2D pictorial structure model of an animal from one video. None of these methods find spatiotemporal correspondences between different instances of a class.

3. System architecture

Our method takes as input a large set of video shots containing instances of an object class (*e.g.* tigers). These input shots are neither temporally segmented nor pre-aligned in any way. The output is a collection of pairwise correspondences between frames. Each correspondence is both temporal, *i.e.* we find correspondences between frames in different shots, and spatial, *i.e.* we recover the transformation mapping points between the two frames (Fig. 1, bottom).

Overview. Fig. 2 shows an overview of our system. The key idea is to first identify pairs of frame sequences from

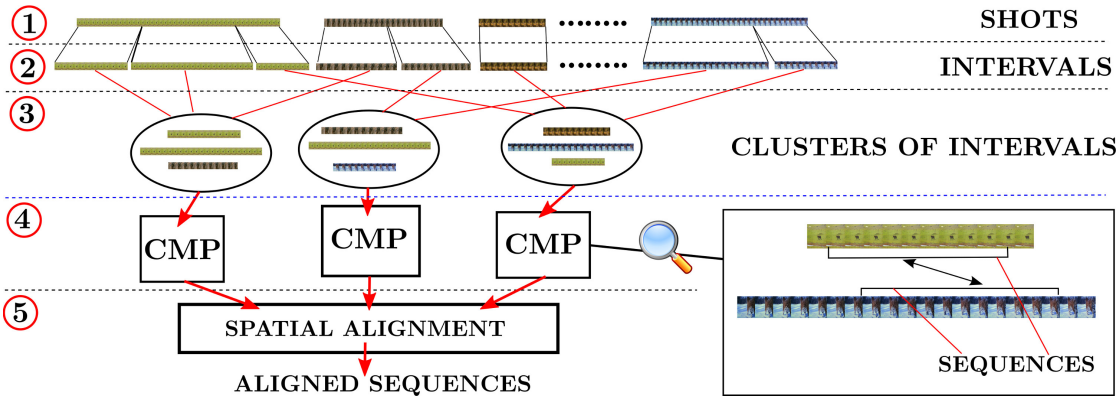

Figure 2. Overview of our method. The input is a collection of shots showing the same class (1). Each shot (which can be of any length) is partitioned into shorter temporal intervals of 10–200 frames (2), which are then clustered together (3) using motion cues. ([11] shows that using intervals shorter than the original shots finds more compact clusters.) The clusters effectively limit the search space: we extract CMPs only from pairs of intervals in the same cluster (4). For each pair, we extract CMPs from all possible pairs of sequences of fixed length (10 frames). An example of a pair is shown in the bottom right (sec. 3 and Fig. 3). Last (5), we align the two sequences of each CMP (see sec. 4 and 5).

two videos that exhibit consistent foreground motion. For this, we use [28] to extract foreground masks from each shot using motion cues and [11] to cluster short intervals with similar foreground motion. Within each cluster, we identify pairs of sequences of fixed length $T = 10$ containing similar foreground motion; we term these *consistent motion pairs* (CMPs). By focusing on similar motion, CMPs provide reliable correspondences even across object instances with very different appearance (such as the white and orange tigers in Fig. 1 or the cub and adult in Fig. 6). These are fed to the next stage, which spatiotemporally aligns the two sequences in each CMP.

Foreground masks. We use the fast video segmentation technique [28] to automatically segment the foregound object from the background. These foreground masks remove confusing features on the background and facilitate the alignment process.

CMP extraction. Attempting to spatially align all possible pairs of sequences would be prohibitively expensive (there are over a billion in just 20 minutes of video). Clustering based on motion with [11] significantly limits the search space. We prune further by considering only the top 10 ranked pairs between two intervals p and q in the same cluster according to the following metric (Fig. 3). We describe each frame using a bag of words (BoW) over the Trajectory Shape and Motion Boundary Histogram descriptors [38] of trajectories starting in that frame. Let d_{ij} be the histogram intersection between the BoWs for frame i in p and frame j in q. The similarity between the T-frame sequence pair starting at i and j is

$$\mathrm{s}\Big([f_i^p,\ldots,f_{i+T-1}^p],[f_j^q,\ldots,f_{j+T-1}^q]\Big) = \sum_{t=0}^{T-1} d_{(i+t)(j+t)}\,. \tag{1}$$

This measure preserves the temporal order of the frames, whereas a BoW aggregated over the whole sequences would not. We found this scheme extracts CMPs that reliably show similar foreground motion and form good candidates for spatial alignment.

Sequence alignment. We have explored a variety of approaches for sequence alignment and report on two representative methods here. The first is a coarse rigid alignment generated by fitting a single homography to foreground trajectory descriptors matched across the two sequences (Sec. 4). The second approach fits a non-rigid TPS mapping to edge points extracted from the foreground regions of each frame. This TPS is allowed to deform smoothly through time through the sequence (Sec. 5). Our experiments confirm that the more flexible model outperforms the rigid alignment (Sec. 6).

4. Rigid sequence alignment

Traditionally, homographies are used to model the mapping between two still images, and are estimated from a set of noisy 2D point correspondences [21]. We consider instead the problem of estimating a homography from trajectories correspondences between two sequences (in a CMP). Below we first review the standard approach for still images, and then present our extensions.

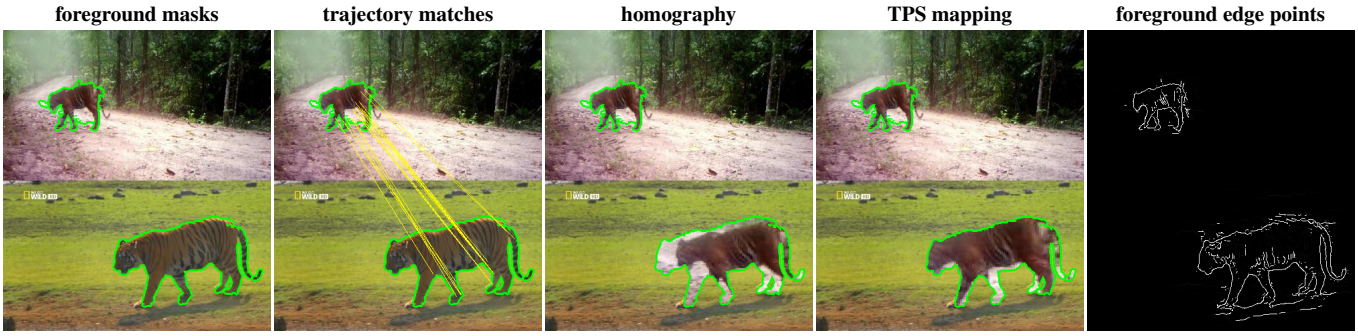

Figure 4. Aligning sequences with similar foreground motion. We first estimate a foreground mask (green) using motion segmentation (a). We then fit a homography to matches between point trajectories (b, sec. 4.2). In (c) we project the foreground pixels in the first sequence (top) onto the second (bottom) with the recovered homography. This global, coarse mapping is often not accurate (note the misaligned legs and head). We refine it by fitting Thin-Plate Splines (TPS) to edge points extracted from the foreground (e, sec. 5). The TPS mapping is non-rigid and provides a more accurate alignment for complex articulated objects (d).

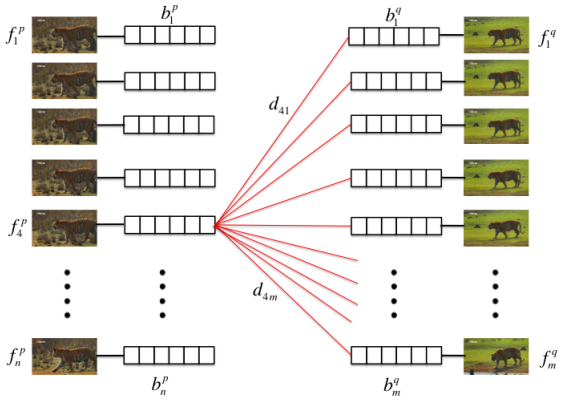

Figure 3. Extracting CMPs from two intervals. First, we approximate the pairwise distance between frames as the histogram distance between their BoWs (which contains all motion descriptors through the frame, sec. 3). Then we keep as CMPs the top scoring pairs of sequences of length T with respect to (1). For the intervals above, the number of pairs of sequences to examine is $(n-T)\cdot(m-T)$.

4.1. Homography between still images

A 2D homography H_{uv} is a 3×3 matrix that can be determined from four or more point correspondences $X_u \leftrightarrow X_v$ by solving

$$X_u = H_{uv} X_v \tag{2}$$

RANSAC [18] estimates a homography from a set of putative correspondences $\mathcal{P}_{uv} = \{(x_u, y_u) \leftrightarrow (x_v, y_v)\}$ that may include outliers. Traditionally, $\mathcal{P}_{uv}$ contains matches between local appearance descriptors, like SIFT [27]. At each iteration, a hypothesis is generated by fitting a homography to four samples from $\mathcal{P}_{uv}$; the computed homography with the smallest number of outliers is kept.

4.2. Homography between video sequences

In video sequences, we use point trajectories as units for matching, instead of SIFT keypoints. We extract trajectories in each sequence and match them using a modified Trajectory Shape (TS) descriptor [38] (Fig. 5). We match each trajectory in the first sequence to its nearest neighbor in the second with respect to Euclidean distance. We use trajectories that are 10 frames long and only match those that start in the same frame in both sequences. Each trajectory match provides 10 point correspondences (one per frame).

We consider two alternative ways to fit a homography to trajectory matches. In the first, we treat the point correspondences generated by a single trajectory match independently during RANSAC. We call this strategy 'Independent Matching' (IM). In the second alternative, we sample four *trajectory* matches at each RANSAC iteration instead of four point correspondences. We solve (2) using the $4T$ associated point correspondences, in the least squares sense. A trajectory match is considered an outlier only if fewer than half of its point correspondences are outliers. We call this strategy 'Temporal Matching' (TM). TM encourages geometric consistency over the duration of the CMP. Instead, IM could overfit to point correspondences from just a few frames. Our experiments show that TM is superior to IM.

4.3. Using the foreground mask as a regularizer

The homography estimated from trajectories tends to be inaccurate when the input matches do not cover the entire foreground (Fig. 6). To address this issue, we note that the bounding boxes of the foreground masks [28] provide a coarse, global mapping (Fig. 7). Specifically, we consider the correspondences between the bounding box cor-

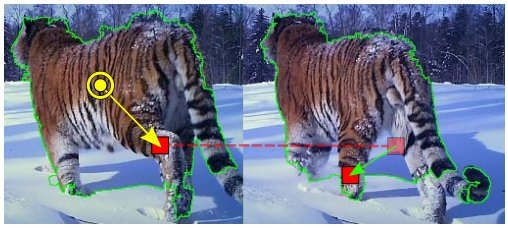

Figure 5. Modifying the TS descriptor. The TS descriptor is the concatenation of the 2D displacement vectors (green) of a trajectory across consecutive frames. This descriptor works well when aggregated in unordered representations like Bag-of-Words [38], but matches found between individual trajectories are not very robust. For example, the TS descriptors for the trajectories on the torso of a tiger walking are almost identical. We make TS more discriminative by appending the vector (yellow) between the trajectory and the center of mass of the foreground mask (green) in the frame where the trajectory starts. We normalize this vector by the diagonal of the bounding box of the foreground mask to preserve scale invariance.

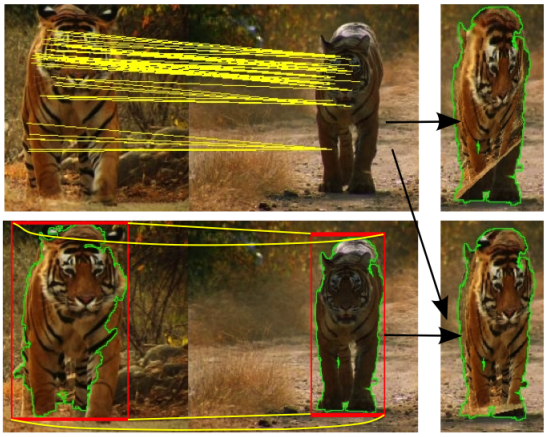

Figure 6. Top: Trajectory matches (yellow) often cover only part of the object (head and right leg here). Here, the homography overfit to correspondences on the head, providing an incorrect mapping for the legs (right). Bottom: Adding correspondences from the foreground bounding boxes provides a more stable mapping (right). The correspondences in the bottom row are also found automatically by our method (no manual intervention needed).

ners, which we call 'foreground matches' (F_u, F_v). These are included in Eq. (2) as additional point correspondences (four per frame):

$$\min \|H_{uv}X_v - X_u\| + \|H_{uv}F_v - F_u\| \quad . \tag{3}$$

This form of regularization makes our method much more stable (Fig. 6).

5. Temporal TPS for sequence alignment

In this section we present a second approach to sequence alignment, based on time-varying thin plate splines

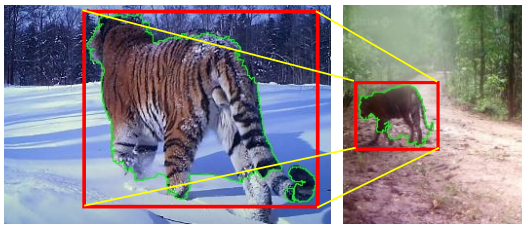

Figure 7. Matching corners between the bounding boxes of the foreground mask provide additional point correspondences between the two sequences. While these correspondences are too coarse to provide a detailed spatial alignment between the sequences, and are sensitive to errors in the foreground segmentation (see Fig. 14), they are useful as a regularizer when combined with other point correspondences.

(TTPS). Unlike the approach presented in the previous section, TTPS is a non-rigid mapping, which is more suitable for putting different object instances in correspondence. We build on the popular TPS Robust Point Matching algorithm [7], originally developed to align sets of points between two still images (Sec. 5.1). We extend TPS-RPM to align two sequences of frames with a TPS that evolves smoothly over time (Sec. 5.2).

5.1. TPS-RPM

A TPS is a smooth, non-rigid mapping, f, comprising an affine transformation d and a non-rigid warp w. The mapping is a single closed-form function for the entire space, with a smoothness term $\mathrm{L}(\mathrm{f})$ defined as the sum of the squares of the second derivatives of f over the space [7]. Given two sets of points $\mathcal{U} = \{u_i\}$ and $\mathcal{V} = \{v_i\}$ in correspondence, f can be estimated by minimizing

$$\mathrm{E}(\mathrm{f}) = \sum_i ||u_i - \mathrm{f}(v_i)||^2 + \lambda ||\mathrm{L}(\mathrm{f})||. \tag{4}$$

$\mathcal{U}$ and $\mathcal{V}$ are typically the position of detected image features (we use edge points, sec. 5.2).

As the point correspondences are typically not known beforehand, TPS-RPM jointly estimates f and a soft-assign correspondence matrix $M = \{m_{ij}\}$ by minimizing

$$\mathrm{E}(M, \mathrm{f}) = \sum_i \sum_j m_{ij} ||u_i - \mathrm{f}(v_j)||^2 + \lambda ||\mathrm{L}(\mathrm{f})||. \tag{5}$$

TPS-RPM alternates between updating f by keeping M fixed, and the converse. M is continuous-valued, allowing the algorithm to evolve through a continuous correspondence space, rather than jumping around in the space of binary matrices (hard correspondence). It is updated by setting m_{ij} as a function of the distance between u_i and $\mathrm{f}(v_j)$ [7]. The TPS is updated by fitting f between $\mathcal{V}$ and the current estimates $\mathcal{Y}$ of the corresponding points, computed

from $\mathcal{U}$ and M. TPS-RPM and optimizes (5) in a deterministic annealing framework, which allows TPS-RPM to find a good solution even when starting from a relatively poor initialization.

5.2. Temporal TPS

Our goal is to find a series of mappings $\mathcal{F} = \{\mathrm{f}^1, \ldots, \mathrm{f}^T\}$, one at each frame in the input sequences. We enforce temporal smoothness by constraining each mapping to use a set of point correspondences that is consistent over time. Let $\mathcal{U}^t = \{u_i^t\}$ be a set of points for frame t in the first sequence (with $\mathcal{V}^t$ defined analogously). This set contains both edge points extracted in t as well as edge points extracted in other frames of the sequence and propagated to t via optical flow (Fig. 9). Each $\mathcal{U}^t$ stores points in the same order such that u_i^t and u_i^τ are related by flow propagation. We solve for $\mathcal{F}$ by minimizing

$$\mathrm{E}(\mathcal{M}, \mathcal{F}) = \sum_t \left(\sum_i \sum_j m_{ij}^t ||u_i^t - \mathrm{f}^t(v_j^t)||^2 + \lambda ||\mathrm{L}(\mathrm{f}^t)|| \right). \tag{6}$$

subject to the constraint that $m_{ij}^t = m_{ij}^\tau$. That is, if two points are in correspondence in frame t, they must still be in correspondence after being propagated to frame τ.

Inference. Minimizing (6) is very challenging. In practice, we find an approximate solution by first using TPS-RPM to fit a separate TPS f^t to the edge points extracted at time t only. This is initialized with the homography found in Sec. 4.3. f^t fixes the correspondences, which we use to estimate f^τ in all other frames. We repeat this process starting in each frame, generating a total of T TTPS candidates and keep the highest scoring one according to (6). Thanks to this efficient approximate inference, we can apply TTPS to align thousands of CMPs.

Foreground edge points. We extract edges using [12]. We remove clutter edges far from the object by multiplying the edge strength of each point with the Distance Transform (DT) of the image with respect to the foreground mask (*i.e.*, the distance of each pixel to the closest point on the mask). We prune points scoring ≤ 0.2. This removes most background edges, and is robust to cases where the mask does not cover the complete object (Fig 8). To accelerate the TTPS fitting process, after pruning we subsample the edge points to at most 1,000 per image.

6. Evaluation

We evaluate our method on 108 shots of tigers from a dataset of documentary nature footage [11] and 98 shots of horses from YouTube-Objects [29], for a total of 17,000 frames per class (roughly 25 minutes of video).

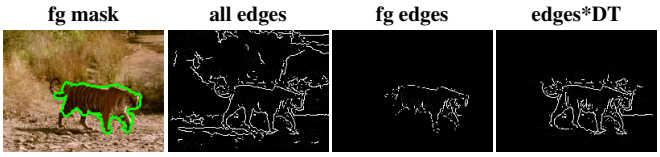

(a) (b) (c) (d)

Figure 8. Edge extraction. Using edges extracted from the entire image confuses the TPS fitting due to background edge points (b). Using only edges on the foreground mask (c) loses useful edge points if the mask is inaccurate, *e.g.* the missing legs in (a). We instead weigh the edge strength (b) by the Distance Transform (DT) with respect to the foreground mask. This is robust to errors in the mask, while pruning most background edges (d).

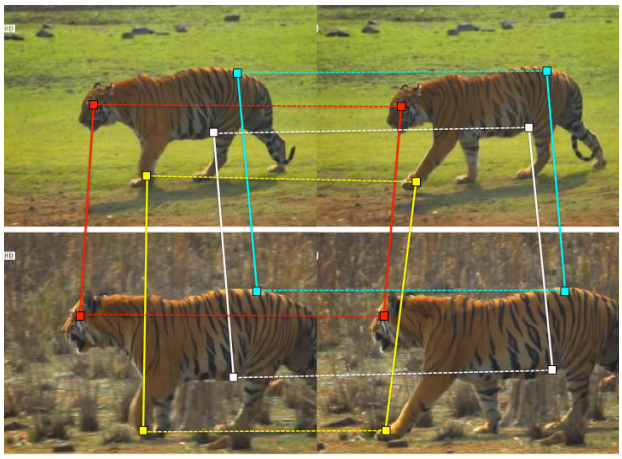

Figure 9. Propagation using optical flow. In each sequence, we propagate edge points extracted at time t using optical flow, independently in each sequence (dashed lines). Our TTPS model (Sec. 5.2) enforces that the correspondences between edge points at time t (solid lines) are consistent with their propagated version at time $t+1$.

6.1. Evaluation protocol

Landmark annotations. In each frame, we annotate the 2D location of 19 landmarks on each tiger/horse[1] (such as eyes, knees, chin, Fig. 10). We do not annotate occluded landmarks. We will make these annotations publicly available. Unlike coarser annotations, such as bounding boxes, landmarks enable evaluating the alignment of objects with non-rigid parts with greater accuracy.

Evaluation measure. We evaluate the mapping found between the two sequences in a CMP as follows. For each frame, we map each landmark in the first sequence onto the second and compute the Euclidean distance to its ground-truth location. The evaluation measure is the average be-

[1] If multiple are visible, we annotate the animal closest to the camera.

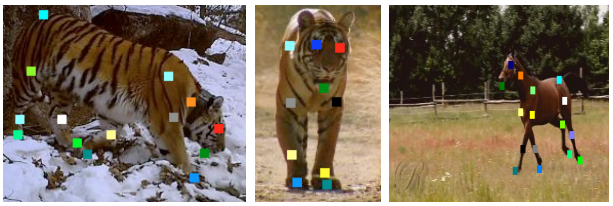

Figure 10. Examples of annotated landmarks. A total of 19 points are marked when visible in over 17,000 frames for two different classes (horses and tigers). Our evaluation measure uses to landmarks to evaluate the quality of a sequence alignments (sec. 6.1).

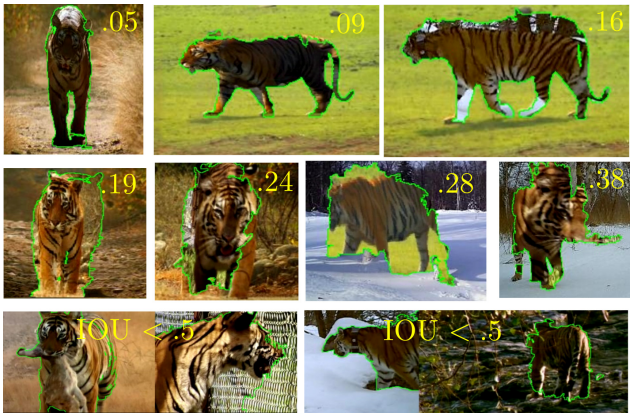

Figure 11. Evaluation measure. We use the ground-truth landmarks to measure the alignment error of the mappings estimated by our method (sec. 6.1). As the error increases, the quality of the alignment clearly degrades. Around 0.18 the alignments contain some slight mistakes (*e.g.*, the slightly misaligned legs in the top right image), but are typically acceptable. We consider a mapping incorrect also when the IOU of the visible landmarks in the aligned pair is below 0.5 (bottom row).

tween this distance and the reverse (*i.e.*, the distance for landmarks mapped from the second sequence into the first). We normalize the error by the scale of the object, defined as the maximum distance between any two landmarks in the frame. The overall error for a pair of sequences is the average error of all visible landmarks over all frames.

After visual inspection of many sampled alignments (Fig. 11), we found that 0.18 was a reasonable threshold for separating acceptable alignments from those with noticeable errors. We count an alignment as *correct* if the error is below this threshold and if the Intersection over Union (IOU) of the two sets of visible landmarks in the sequence is above 0.5 (to avoid rewarding accidental alignments of a few landmarks, bottom row of fig. 11).

6.2. Evaluating CMP extraction

First, we evaluate our method for CMP extraction in isolation (sec. 3). Given a CMP, we use the ground-truth landmarks to fit a homography, and check if it is correct according to the evaluation measure above. If so, it means that it is in principle possible to align it (we call it *alignable*). Our method returns roughly 3000 CMP on the tiger data, of which 51% are alignable. As a baseline, we consider extracting CMPs by uniformly sampling sequences from pairs of shots. In this case, the percentage of alignable CMPs drops to 19%. Results are similar on the the horse dataset: our method delivers 49% alignable CMPs, vs 26% by the baseline.

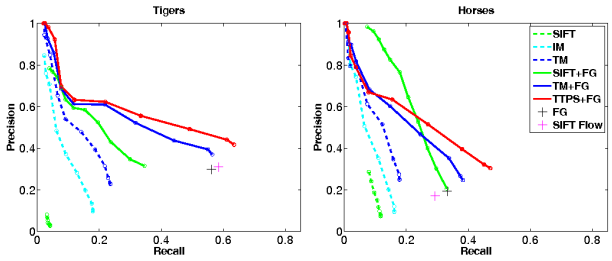

Figure 12. Evaluation of sequence alignment. We separately evaluate our method on two classes, horses and tigers. With no regularization, trajectory methods are superior to SIFT on both classes, with TM performing better than IM. Adding regularization using the foreground matches improves the performance of both TM and SIFT (compare the dashed to the solid curves). TTPS clearly outperform all trajectory methods, as well as SIFT Flow and the FG baseline (see text).

6.3. Evaluating spatial alignment

We now evaluate various methods for automatic sequence alignment. For each method, we generate a precision-recall curve as follows. Let n be the total number of CMPs returned by the method; c the number of correctly aligned CMPs; and a the total number of alignable CMPs (sec. 6.2). Recall is c/a, and precision is c/n. Different operating points on the precision-recall curve are obtained by varying the maximum percentage of outliers allowed when fitting a homography.

Comparison to other methods. We compare our method against SIFT Flow [26]. We use [26] to align each pair of frames from the two sequences independently. We restrict the algorithm to match only the bounding boxes of the foreground masks, after rescaling them to be the same size (without these two steps, performances significantly drop).

Further, we also compare to fitting a homography to SIFT matches found in the two sequences. We use only keypoints on the foreground mask, and preserve temporal order by matching only keypoints in corresponding frames. We tested this method alone (SIFT), and by adding spatial regularization using the foreground masks (SIFT + FG, sec. 4.3).

Finally, we consider a simple baseline that fits a homography to the bounding box of the foreground masks alone

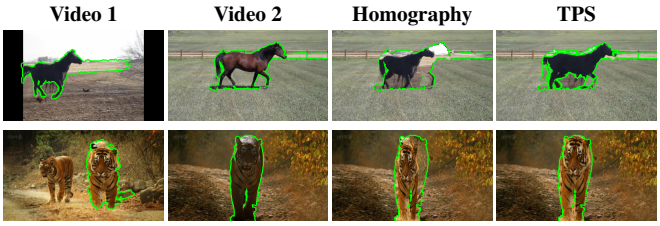

Figure 13. TTPS (d) provide a more accurate alignment for complex articulated objects than homographies (c).

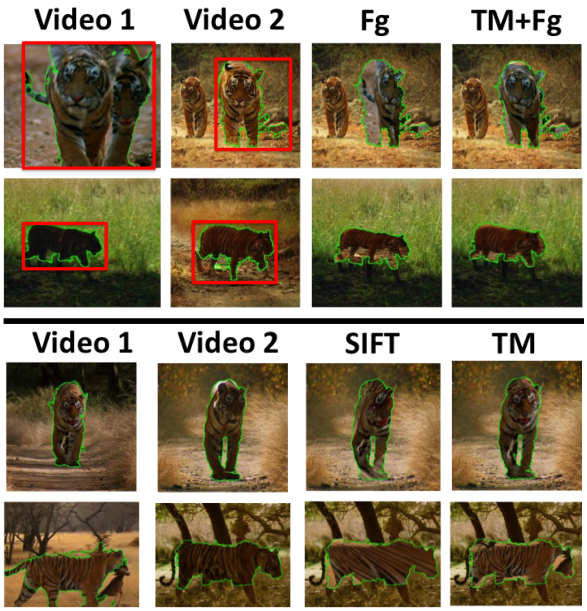

Figure 14. Top two rows: Estimating the homography from the foreground masks alone fails when the bounding boxes are not tight around the objects (first-second columns). Adding trajectories (TM+FG) is more accurate (fourth column, sec. 4.2). Bottom two rows: the striped texture of tigers often confuses estimating the homography from SIFT keypoint matches (third row). On this class, using trajectories (TM) often performs better.

(FG).

Analysis of rigid alignment. Both trajectory methods (TM, IM, sec. 4.2) are superior to SIFT on both classes, with TM performing better than IM (Fig. 12). Adding spatial regularization with the foreground masks (+FG) improves the performance of both TM and SIFT. SIFT performs poorly on tigers, since the striped texture confuses matching SIFT keypoints (Fig. 14, bottom). Trajectory methods work somewhat better on tigers than horses due to the poorer quality of YouTube video (*e.g.* low resolution, shaky camera, abrupt pans). As a result of these factors, TM+FG clearly outperforms SIFT+FG on tigers, but it is somewhat worse on horses.

Analysis of TTPS. The time-varying TPS model (TTPS+FG, sec. 5) significantly improves upon its initialization (TM+FG) on both classes. On tigers, it is the best method overall, as its precision-recall is above all other curves for the entire range. On horses, the SIFT+FG and TTPS+FG curves intersect. However, TTPS+FG achieves a higher Average Precision (*i.e.* the area under the curve): 0.265 vs 0.235.

The SIFT Flow software [26] does not produce scores comparable across CMPs, so we cannot produce a full precision-recall curve. At the level of recall of SIFT Flow, TTPS achieves +0.2 higher precision on tigers, and +0.3 on horses. We also note that TM and TM+FG are closely related to the method for fitting homographies to trajectories in [6]. TM+FG augments [6] in several ways (automatic CMP extraction, modified TS desriptor, regularization with the foreground masks), but is still inferior to TTPS+FG. Last, TTPS also achieves a significantly higher precision than the FG baseline. This shows that our method is robust to errors in the foreground masks. In supplemental material we provide example head-to-head qualitative results, showing that TTPS alignents typically look more accurate than the other methods (Fig. 13).

For the tiger class, out of all CPMs returned by TTPS (rightmost point on the curve), $1,000$ of them are correctly aligned tiger (*i.e.* $10,000$ frames). The precision at this point is 0.5, i.e. half of the returned CMPs are correctly aligned. For the horse class, TTPS returns 800 correctly aligned CMPs, with precision 0.35.

7. Discussion

We present a method that automatically extracts dense spatiotemporal correspondences from a collection of videos showing a particular object class. Our pipeline consumes raw video, without the need for manual annotations or temporal segmentation. Using motion as the primary signal for identifying correspondences allows us to match sequences despite significant appearance variation. Ultimately, the thin plate spline matching results in temporally-stable, high-quality alignments for thousands of sequence pairs.

Our method is not limited to a particular class of object but instead applies to any objects that exhibit consistency in behavior and thus exhibit the same characteristic motion patterns across different observations. The correspondences we find can be used to learn a general model of the object class without requiring any human supervision beyond video-level object class labels. Additionally, they can enable novel applications, such as replacing an instance of an object with an instance from a different video, or retrieving videos in a collection that tightly match the motion of the object in a query video.

Acknowledgments. We are very grateful to Anestis Papazoglou for helping with the data collection, and to Shumeet Baluja for his helpful comments. This work was partly funded by a Google Faculty Research Award, and by ERC Starting Grant "Visual Culture for Image Understanding".

References

[1] H. Azizpour and I. Laptev. Object detection using strongly-supervised deformable part models. In *ECCV*, 2012. 1

[2] C. Barnes, E. Shechtman, D. Goldman, and A. Finkelstein. The generalized patchmatch correspondence algorithm. In *ECCV*, 2010. 1, 2

[3] L. Bourdev and J. Malik. Poselets: Body part detectors trained using 3d human pose annotations. In *ICCV*, 2009. 1

[4] M. Brown and D. Lowe. Automatic panoramic image stitching using invariant features. *IJCV*, 74(1), 2007. 2

[5] Y. Caspi and M. Irani. A step towards sequence-to-sequence alignment. In *CVPR*, 2000. 2

[6] Y. Caspi, D. Simakov, and M. Irani. Feature-based sequence-to-sequence matching. *IJCV*, 68(1):53–64, 2006. 2, 8

[7] H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *CVIU*, 89(2-3):114–141, Feb 2003. 2, 5

[8] O. Chum and J. Matas. Optimal randomized ransac. *IEEE Trans. on PAMI*, 2008. 2

[9] R. Cinbis, J. Verbeek, and C. Schmid. Segmentation driven object detection with fisher vectors. In *ICCV*, 2013. 1

[10] N. Dalal and B. Triggs. Histogram of Oriented Gradients for human detection. In *CVPR*, 2005. 1

[11] L. Del Pero, S. Ricco, R. Sukthankar, and V. Ferrari. Articulated motion discovery using pairs of trajectories. In *CVPR*, 2015. 2, 3, 6

[12] P. Dollar and C. Zitnick. Structured forests for fast edge detection. In *ICCV*, 2013. 6

[13] Q. Fan, K. Barnard, A. Amir, and A. Efrat. Robust spatio-temporal matching of electronic slides to presentation videos. *IEEE Transactions on Image Processing*, 20(8):2315–2328, 2011. 2

[14] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *IEEE Trans. on PAMI*, 32(9), 2010. 1

[15] P. F. Felzenszwalb and D. P. Huttenlocher. Pictorial structures for object recognition. *IJCV*, 61(1):55–79, 2005. 1

[16] V. Ferrari, F. Jurie, and C. Schmid. From images to shape models for object detection. *IJCV*, 87(3), 2010. 2

[17] V. Ferrari, T. Tuytelaars, and L. Van Gool. Simultaneous object recognition and segmentation from single or multiple model views. *IJCV*, 67(2):159–188, 2006. 2

[18] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. ACM*, 24(6):381–395, 1981. 2, 4

[19] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 1

[20] L. Gorelick, M. Blank, E. Shechtman, M. Irani, and R. Basri. Actions as space-time shapes. *IEEE Trans. on PAMI*, 29(12):2247–2253, December 2007. 2

[21] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000. 1, 3

[22] A. Jain, A. Gupta, M. Rodriguez, and L. Davis. Representing videos using mid-level discriminative patches. In *CVPR*, 2013. 2

[23] H. Jegou, M. Douze, and C. Schmid. Hamming embedding and weak geometric consistency for large-scale image search. In *ECCV*, 2008. 2

[24] Y. Ke, R. Sukthankar, and M. Hebert. Event detection in crowded videos. In *ICCV*, 2007. 2

[25] J. Liao, R. S. Lima, D. Nehab, H. Hoppe, and P. V. Sander. Semi-automated video morphing. In *Eurographics Symposium on Rendering*, 2014. 2

[26] C. Liu, J. Yuen, A. Torralba, J. Sivic, and W. Freeman. SIFT Flow: Dense correspondence across different scenes. In *ECCV*, 2008. 1, 2, 7, 8

[27] D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004. 1, 2, 4

[28] A. Papazoglou and V. Ferrari. Fast object segmentation in unconstrained video. In *ICCV*, December 2013. 3, 4

[29] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Learning object class detectors from weakly annotated video. In *CVPR*, 2012. 2, 6

[30] D. Ramanan, A. Forsyth, and K. Barnard. Building models of animals from video. *IEEE Trans. on PAMI*, 28(8):1319 – 1334, 2006. 2

[31] C. Rao, A. Gritai, and M. Shah. View-invariant alignment and matching of video sequences. In *ICCV*, 2003. 2

[32] C. Schmid and R. Mohr. Combining greyvalue invariants with local constraints for object recognition. Technical report, INRIA Rhône-Alpes, Grenoble, France, 1996. 2

[33] S. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *CVPR*, 2006. 2

[34] K. Tang, R. Sukthankar, J. Yagnik, and L. Fei-Fei. Discriminative segment annotation in weakly labeled video. In *CVPR*, 2013. 2

[35] Y. Ukrainitz and M. Irani. Aligning sequences and actions by maximizing space-time correlations. In *ECCV*, 2006. 2

[36] P. A. Viola, J. Platt, and C. Zhang. Multiple instance boosting for object detection. In *NIPS*, 2005. 1

[37] G. Wahba. *Spline models for observational data*, volume 59 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, 1990. 2

[38] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013. 3, 4, 5

[39] X. Wang, M. Yang, S. Zhu, and Y. Lin. Regionlets for generic object detection. In *ICCV*, pages 17–24. IEEE, 2013. 1

[40] A. Yilmaz and M. Shah. Actions as objects: A novel action representation. In *CVPR*, 2005. 2